

Combining the ApEn statistic with surrogate data analysis for the detection of nonlinear dynamics in time series

Randall A. LaViolette,^{*} Charles R. Tolle,
 Timothy R. McJunkin and Daphne L. Stoner
 IDAHO NATIONAL ENGINEERING & ENVIRONMENTAL LABORATORY[†]
 P.O. BOX 1625, IDAHO FALLS ID, 83415, USA

February 8, 2008

Abstract

We tested the natural combination of surrogate data analysis with the ApEn regularity statistic developed by Pincus [*Proc. Natl. Acad. Sci. USA* **88** (1991) 2297] by applying it to some popular models of nonlinear dynamics and publicly available experimental time series. We found that this easily implemented combination provided a useful method for discriminating signals governed by nonlinear dynamics from those governed by linear dynamics and noise. An apparently novel physical interpretation of ApEn also is supplied.

1 Introduction

Significant developments in the last decade have provided techniques that can discriminate between signals generated by essentially linear dynamics and signals generated by essentially nonlinear and possibly chaotic dynamics[1][2][3][4][5].

^{*}Mailstop: 2208, e-mail: yaq@inel.gov

[†]Operated for the U. S. Department of Energy under DOE-ID Operations Office Contract DE-AC07-99ID13727.

Our goal here is to assess the utility of the combination of surrogate data techniques with the “ApEn” (or “approximate entropy”) statistic for the identification of essentially nonlinear signals. In the next two subsections we briefly recapitulate surrogate data methodology and the ApEn statistic, respectively. We present the results of applying the combined analysis to various data sets and models of nonlinear dynamics in the next Section. We conclude in Section 3 with some comments on the results. Three short appendices follow that set forth some details of the algorithm for calculating ApEn, the connection between ApEn and standard quantities, and the previously unnoticed numerical coincidence of ApEn with the thermodynamic entropy of the one-dimensional Ising model, respectively.

1.1 Surrogate Data Methodology

The idea of surrogate data is that one tests some hypothesis concerning a time series by comparing the value of some statistic Q for the original time series with the sample mean of that statistic \bar{Q} in an ensemble of random replicates (the surrogates) generated from the original time series under that hypothesis. The hypothesis would be taken to be true if the original time series were statistically indistinguishable from the ensemble of surrogate time series. “The use of surrogate data is essential for deciding whether an irregular time series arises from nonlinear deterministic chaos or linear stochastic dynamics. The method was introduced by Theiler et al. (1992).” [3] The particular method that we are concerned with was designated FT by Theiler et al. [6][7][8], and tests the following linearity hypothesis: The original time series is the result of linear stationary dynamics driven by gaussian white noise, i.e., the data can be modeled by ARMA dynamics[5]. The surrogates are produced by first randomizing the phase of the Fourier transform of the original time series, then inverting back to the time domain to create a new time series in the ensemble of surrogate time series. This procedure ensures that each surrogate time series possesses the same power spectrum as the original data. Theiler et al. suggested that the “significance”, i.e., the number S of standard deviations $\sigma_{\bar{Q}}$ of \bar{Q} between Q and \bar{Q} , written as

$$S = \frac{|Q - \bar{Q}|}{\sigma_{\bar{Q}}}, \quad (1)$$

could be employed as a simple albeit crude test for the success or violation of the linearity hypothesis. They also suggested that the statistic Q could be

the correlation integral or correlation dimension, suggestions that have been widely adopted, although other statistics also have been suggested for the surrogate data method[9][10][11][12][13]. In view of the usual uncertainties in the estimate of $\sigma_{\bar{Q}}$ and the crudeness of S itself[6], in this work we required $S \geq 10$ to score a violation, although other authors[13] have been willing to declare a violation of the linearity hypothesis with a significance as small as three. The analysis is meaningful only for stationary time series; examples of violations of the linearity hypothesis by linear non-stationary processes have been documented[14].

Several authors have proposed modifications to the FT method for generating surrogate series. Theiler et al. themselves introduced the “amplitude-adjusted” method (designated AAFT by them[6]) because the amplitudes of the surrogate data generated by the FT method can vary significantly from the envelope of the original time series. It turned out that the AAFT itself introduced other concerns that in turn have been addressed by yet more elaborate and in some cases more computationally demanding surrogate data methods[5][15][16]. Those proposals to modify the FT method retained the original suggestions for the statistics; here we pursue the opposite approach, retaining FT here for computational and conceptual simplicity, and employing ApEn as an alternative statistic. Of course ApEn could have been just as easily employed with any of the other methods to generate a surrogate series.

1.2 ApEn: A regularity statistic

In this paper we address only the choice of the statistic in the FT method. One of our concerns with both the correlation integral and the correlation dimension as a statistic is that each may require a large number of data points. The data required may grow rapidly with the correlation dimension, so that some workers have had to employ this or related statistics in a regime where the validity of doing so was difficult to assess[17][18]. One of the appeals of ApEn is that it has performed well, in other contexts[19][20][21][22][23][24], even on decimal data sets as small as $N \approx 100$, or binary data sets as small as $N \approx 20$. Therefore we investigated the application of ApEn to the FT surrogate data method, which we designate ApEn+FT.

The statistic ApEn was developed by Pincus to measure the conditional probability that “...runs of patterns that are close for [some number] of observations remain close...”[19]. ApEn therefore depends upon the length (N)

of the series, the width (m) of the window that defines the patterns, and the tolerance (r) that defines closeness of the patterns, and incidentally provides ApEn with some resistance to noise. Pincus' algorithm for the calculation of ApEn is recapitulated in the Appendix A for reference. The theoretical (large N , m and small r) bounds to ApEn are zero for a perfectly regular sequence, and $\ln(B)$ for a maximally irregular (i.e., random) sequence of base B numbers. It turns out (see Appendix C) that ApEn numerically coincides with the thermodynamic entropy for the one-dimensional Ising model when each configuration of spins is interpreted as a string of bits. This coincidence serves at least to provide a physical interpretation of ApEn (for binary data), even though ApEn was constructed without any explicit reference to statistical mechanics or thermodynamics. However, we did not further exploit this connection. Furthermore, in spite of its name and its asymptotic behavior (see Appendix B), we did not regard ApEn as an approximation to the, e.g., Kolmogorov-Sinai or Eckmann-Ruelle, non-thermodynamic entropies because the amount of data required to achieve useful approximations may be very large in some circumstances[17][18][19], and because those entropies can be especially sensitive to noise. We also did not compare ApEn with either the correlation integral or correlation dimension for three reasons. First, we have not attempted to establish the superiority of ApEn over all other statistics in all circumstances. Second, both the correlation integral and the correlation dimension nominally require the evaluation of limits that in turn may also require much larger quantities of data than ApEn itself[19] (see also Appendix B). Third, the comparison between ApEn and the correlation integral (for example) is complicated by the fact that the choice of parameters m and r that are optimal for one are almost certainly not optimal for the other; employing only the parameters that we also employed for the ApEn calculation might have appeared to be creating a "straw man".

Therefore, we adopted the alternative recommended by Pincus[21], i.e., that ApEn is simply a regularity statistic (or family of statistics) defined by the choice of its (N, m, r) parameters that stands on its own. As such, ApEn has been applied to binary sequences[23][24] and to physiological data[25][26][27]. Although ApEn is a biased statistic, as are many nonlinear statistics, we always employed it to compare only the sequences with the same set of (N, m, r) parameters, so that the bias would not become an issue here. ApEn is strongly sensitive to the complexity of the data, but is insensitive to other attributes, as would be desired of such a statistic. For example, and not surprisingly, ApEn is insensitive to topological conjugation[2] of sequences.

Furthermore, ApEn is semi-pivotal in the sense that it is always insensitive to the sample mean of the data, and it is insensitive to the sample variance if, as Pincus suggests[19], the tolerance parameter r is always chosen as a fixed fraction of the sample standard deviation. This feature contributed to the utility of ApEn in surrogate data methodology[8]; in particular, the concern that inspired the AAFT surrogate data method should be insignificant here.

Even though the combination of ApEn with surrogate data methodology is a natural one to employ, we found no examples of such analysis in the physical science or engineering literature. However we recently discovered two examples of this kind of analysis in the cardiology literature[26][27], although those applications appeared to us to be narrow, with no attempt to test the method with standard data or models, or to explore its limitations. Therefore we proceed to show the results of our tests of ApEn+FT in the next section.

2 Results of tests with ApEn+FT

The plan of this section consists of the employment and empirical testing of the ApEn statistic in the FT surrogate data method (ApEn+FT) to both experimental and computational data. In order to obtain consistent results, we followed published recommendations[19][21] to fix the pattern window length to $m = 2$ and the tolerance to $r = 0.2\sigma$, where σ was the sample standard deviation estimated by the non-parametric bootstrap[28]. Although important questions remain concerning the optimal values of N , m , and r , we make no pretense of answering them here but defer them to future work. However, preliminary estimates (for decimal numbers) suggest that the minimum number of points for an adequate calculation of ApEn requires on average only about 30^m points[21]. Therefore we always considered data sets at least as large as $N = 2048$. On the other hand, we did not examine data sets larger than $N = 8192$ ($N = 4096$ for most data sets), because we wanted to see if the ApEn+FT analysis would be useful for medium-size data sets; for very large data sets it is conceivable that any number of methods would work equally well, but such large data sets are rarely achieved experimentally. We also restricted N to powers of two so that we could use the simplest version of the Fast Fourier Transform[29] without having to pad the data.

For each original time series we generated an ensemble of only ten FT

surrogate time series, because the fluctuation of ApEn in the ensemble was consistently small and because occasional tests with 40 surrogates did not yield significantly different estimates. As discussed below, the length N of the series had a much greater impact on the performance of the hypothesis test than did the size of the surrogate ensemble. The percentile bootstrap method was employed to estimate the (0.5%,99.5%) or (0.1%,99.9%) confidence intervals employed in the Figures in order to assist visual inspection, but only the bootstrap estimate of the standard deviation was employed in the significance calculations (see Equation 1)[28]. The bootstrap usually provides a good estimate of the sample standard deviation even for ten samples, especially where, as was the case here, the median and the mean are nearly coincident.

We did not display the time series themselves because we found that it was misleading to visually compare the original time series with its surrogates. The surrogate data do not need to overlap point for point with the original in order for them to be statistically the same. The FT method randomizes the phase, but ApEn is properly insensitive to the phase (if any) present in the data, so that even when the original and surrogates have the same ApEn, the time series themselves may appear to be quite different from one another. On the other hand, some series (e.g., Hénon) vary so rapidly, that the original and surrogates may appear similar when they are actually statistically distinct.

2.1 Noise

We began by studying the behavior of ApEn for white noise, in order to provide a context for the values of ApEn for the subsequent data, and to motivate the combination of ApEn with the surrogate data techniques. We generated an ensemble of ten time series from a white noise source provided here by two uniform pseudorandom number generators, respectively: a standard multiplicative-congruential generator, implemented as “ran2” [29], and a much different generator employing a random-walk algorithm implemented as “rawkrab” [30]. Each generator gave the same result for the sample mean, i.e., $\text{ApEn}(N = 4096, m = 2, r = 0.2\sigma) = 2.11$ (with a standard deviation estimated to be 0.01), about 9% below the theoretical maximum $\text{ApEn}(N \rightarrow \infty, m \rightarrow \infty, r \rightarrow 0) = \ln(10)$ [19]. We continued this calibration by generating another ensemble of ten time series directly from the Brownian

noise (or Ornstein-Uhlenbeck) power spectrum[31]

$$S_0(\omega, \gamma) = \frac{2\gamma\sqrt{N}}{\pi(\omega^2 + \gamma^2)}, \quad (2)$$

for each dissipation rate γ , i.e., we assigned random phases to S_0 and computed its inverse Fourier transform to obtain each trajectory[32]. Such trajectories automatically satisfy the linearity hypothesis. The resulting time series is a solution of a Langevin equation, which can be represented simply for discrete times by the ARMA model[33] $x_n = c_0 + c_1x_{n-1} + c_2\xi_n$, with time-invariant coefficients $c_j(N, \gamma)$, and the uncorrelated gaussian noise of zero mean and unit variance ξ_n . Figure 1 shows the sample mean ApEn computed for each of the various γ ; the standard deviation (estimated by bootstrap techniques[28]) is about 0.01, or about 1/10 the diameter of the circles. For small γ , the spectrum becomes a smeared delta function, the trajectories become more regular, and ApEn is near zero. For larger γ , the spectrum becomes flatter, the trajectories become more random, and ApEn approaches its white noise limit (i.e., 2.11) for the parameters employed above. For intermediate values of γ the trajectories appear noisy but more regular, as shown in the inset of Figure 1 for $\gamma = 1/64$. Thus any allowed value for ApEn can be found from the Brownian time series. Therefore, instead of relying upon the value of ApEn alone, we combined it with surrogate data generation in order to test the hypothesis of linear dynamics.

2.2 Data sets and models

We applied ApEn+FT to data sets and models that have been widely discussed and for which the time series can be either readily computed or freely downloaded, as follows: the data sets labeled “A” (chaotic laser), “D” (turbulent flow), and “E” (light variability from a star) were taken from the Santa Fe Institute’s 1991-1992 competition[34][35] (but not the other data sets that were represented as significantly non-stationary); an experimental realization of the Chua circuit[36][37]; the x-coordinate of the Lorenz system[4]; the x-coordinate of the Hénon map[2]; the Mackey-Glass model[3]; and the x-coordinate of the two-dimensional map of the non-chaotic strange attractor of Grebogi, Ott, Pelikan, and Yorke[38]. In all we obtained 23 original time series, for which we first computed ApEn, as shown in Figure 2. For each of these series we also generated an ensemble of ten FT surrogate time series,




Figure 1: ApEn for white and colored noise. Each circle corresponds to sample mean of ApEn computed for an ensemble of ten trajectories generated (via uniformly random phases) from the Brownian noise spectrum $S_0(\omega, \gamma)$ (Equation(2)), for various values of γ . The standard deviation is estimated to be 0.01, about 1/10 the diameter of the circles. The inset shows a typical time series reconstructed from $S_0(\omega, \gamma = 1/64)$. For reference, the dashed line shows the value of ApEn uniformly distributed (pseudo)random numbers on the interval $[0, 1]$. The dotted line shows the value of $\ln(10)$, the ApEn for uniform random decimal numbers in the limit of large N and m and small r . The circle at the origin is theoretical, not computed.

and calculated the sample-mean ApEn for each ensemble. We calculated (0.5%, 99.5%) confidence intervals about by the non-parametric percentile bootstrap method[28][39]. The results are displayed in Figure 2 and Table 1.

Table 1: Significance of ApEn+FT test

Data Set	N	S
A (mean)	4096	120
D (mean)	4096	50
E (all)	2048	< 5
chua	4096	30
lorenz	8192	300
henon	4096	200
mackey-glass	4096	5
mackey-glass	8192	25
strange	4096	230

Both visual inspection of Figure 2, and the estimation (see Equation 1) of the significance S (Table 1) indicated that the series labeled *A*, *D*, *chua*, *lorenz*, *henon*, *mackey-glass*, and *strange* clearly violate the linearity hypothesis, as expected. To put these results into perspective, others[13] have been willing to declare a violation of the linearity hypothesis with a significance as small as 3 (using other statistics), while we insisted upon a significance of at least 10 (see Equation 1). The small magnitude of the fluctuations of ApEn for the surrogate data contributes substantially to the large values of S in the case of violations of the linearity hypothesis. The Mackey-Glass system is the smoothest and most nearly periodic of any of the time series that we considered here and was also one of the most resistant to our analysis. The time series of 4096 points (not shown in Figure 2) resulted in a significance well below our criterion of $S \geq 10$. In order to find a more significant violation of the linearity hypothesis, we generated the 8192-point time series (see Table 1). Again, the fluctuations in the ApEn of the surrogate data were small. On the other hand, the Lorenz system did not need $N = 8192$ for the violation to become apparent with this analysis; we presented the results in order to indicate the extreme significance that could be achieved with this analysis.




Figure 2: ApEn of original and surrogate time series for various data sets and models. The shading of the columns is used here to group data sets. The top of each column shows the ApEn of the original time series. The height of each pair of error bars shows the (0.5%, 99.5%) bootstrap confidence interval about sample mean (not shown) of the ApEn of the surrogate time series. The labels “A”, “D”, and “E” correspond to the Santa Fe Institute 1991-1992 Competition Data Sets A:continuation, D and E, respectively. Multiple columns labeled A and D correspond to subsets split in increments of 4096 points from the original data set. The columns labeled E correspond to parts 8, 11, 13, and 14, respectively, of Data Set E, and contain only the first 2048 measurements of each part. The column labeled “chua” corresponds to the first 4096 points of a relatively noiseless 5000-point time series of the voltage from an experimentally realized network of two Chua circuits. The next two columns correspond to $N = 8192$ time series generated numerically by us for the x-coordinate of the Lorenz system ($Pr = 16$, $Ra = 45.92$, and $b = 6$) and the x-coordinate of the $N = 4096$ Hénon map ($a = 1.4$ and $b = 0.3$), respectively. The next column corresponds to the $N = 8192$ time series generated numerically by us for the Mackey-Glass model. The last column corresponds to the $N = 4096$ time series generated numerically by us for the x-coordinate of the two-dimensional non-chaotic strange attractor ($\lambda = 1.5$) of Grebogi et al.

The ApEn+FT analysis indicated that series E (light variability of a star) did satisfy the linearity hypothesis, as suggested by both visual inspection of Figure 2 and estimation of the significance (Table 1). If E were actually governed by nonlinear dynamics, then this would have been a failure of our analysis. Our difficulty with the Mackey-Glass system also suggests that our analysis might have failed, because at $N = 2048$, the series in E are the shortest that we have considered. On the other hand, inspection of the data sets themselves and their power spectra suggested that the fluctuations in E might be more random than recent and more substantiated examples of nonlinear light variability of astronomical objects[40][41]. Therefore we were unable to decide between the conclusion that the data in E really are nonlinear but are obscured (to our analysis) by noise, or that the data in E really are governed by ARMA dynamics, as indicated by our analysis.

2.3 Sensitivity to noise

In order to gain a sense of the magnitude of the noise needed to disrupt our analysis, we added uncorrelated gaussian noise $\xi(t)$ (zero mean, unit variance), scaled by a factor, to the experimental Chua circuit and the Hénon map. We chose these time series for comparison because the Chua system violated the linearity hypothesis with relatively small significance ($S \approx 30$) while the Hénon map had a much larger significance ($S \approx 200$). By trial and error we estimated the scale factor ϕ such that the addition of $\phi\xi(t)$ to the time series reduced S to ten. In Table 2 we show the range (i.e., distance between the extreme values) of the original time series, ϕ , and ρ , the corresponding signal-to-noise ratio[29]:

$$\rho = \frac{\sum_t (signal(t))^2}{\sum_t (noise(t))^2}. \quad (3)$$

Table 2: Magnitude ϕ of the noise required to reduce the significance ApEn+FT test to ten

Data Set	Range	ϕ	ρ
chua	6	0.35	40
henon	3	0.45	2.4

The ApEn+FT analysis of the experimental Chua circuit began to fail with the addition of as little as $2\frac{1}{2}\%$ noise, while the artificial Hénon map could withstand up to 42% noise. However, we did not find enough comparable discussions of the signal-to-noise to make comparisons to other methods, so we could not decide from results like these whether our method was superior or inferior to other methods with respect to noise sensitivity.

2.4 Hénon map subject to a linear filter

Experimental data is sometimes filtered even before it becomes “raw” data[42]. Two examples were considered for evenly-sampled discrete data x_n : the moving average (MA) filter in Equation 4, and the linear autoregressive (AR) filter in Equation 5.

$$\bar{x}_n = \frac{1}{2q+1} \sum_{m=-q}^q x_{m+n} \quad (4)$$

$$\hat{x}_n = a \hat{x}_{n-1} + x_n \quad (5)$$

We wanted to assess the impact of these filters on data that were clearly governed by nonlinear dynamics. Therefore we needed data that appeared noisy or jerky because the filters would have little effect on smooth data, (e.g., the Mackey-Glass model) even if they were nonlinear. We chose the data from the Hénon map as described in the subsection above, because of its apparent susceptibility to filtering, and applied both filters, for various values of the filter parameters q and a , respectively. The results are listed in Table 3. The original series (see Figure 2) has an ApEn of 0.48. The sample mean ApEn of the surrogates of the filtered data differed very little from that of the surrogates of the unfiltered data. However, the results in Table 3 indicate that filtering increased the ApEn of the original data by as much as 50%, but the linearity hypothesis still was significantly violated ($S \geq 100$).

2.5 Hénon map tested with AAFT surrogate data

We have said throughout that we did not regard as crucial the particular form of the method employed to generate the surrogate time series. Nevertheless, we observed that the amplitudes of the FT surrogate data for the Hénon map were nearly twice that of the original data (in each direction), the most we observed for any of the data sets, so we thought it prudent to

Table 3: ApEn for MA- and AR-filtered Hénon map. ApEn = 0.48 for the unfiltered Hénon time series.

q (MA)	ApEn	a (AR)	ApEn
3	0.69	0.1	0.46
7	0.76	0.3	0.43
15	0.65	0.5	0.51
31	0.59	0.7	0.69
63	0.58	0.9	0.62

check the consequences of such a difference in the amplitudes by applying the amplitude-adjusted (AAFT) method[6] to this data. Although the AAFT surrogates look much more like the original data, the corresponding sample mean ApEn was found to be 1.94, only about 5% lower than the sample mean ApEn (2.03) of the FT surrogate data; the sample standard deviations, estimated to be 0.006 and 0.008 for the AAFT and the FT surrogates, respectively, also differed little from one another.

2.6 Logistic map with varying parameter

The results above show the behavior of the ApEn statistic with FT surrogate data analysis for a wide range of real data and models. However, none of these provided insight into the behavior of the analysis as the regularity of the data varies. Here we considered the logistic map in order to examine the behavior of the analysis with a parameter that changes the regularity of the data.

The logistic map is usually defined by[1]

$$y_{n+1} = \mu y_n(1 - y_n) \quad (6)$$

or equivalently, with a different choice of coordinates,[2]

$$\begin{aligned} x_{n+1} &= 1 - \rho x_n^2 \\ \rho &= \frac{\mu}{2} \left(\frac{\mu}{2} - 1 \right) \end{aligned} \quad (7)$$

Figure 3 shows the value of ApEn of the time series generated from Equation 6 and its surrogates for various values of μ . For $\mu \leq 3.55$, the dynamics are consistent with the linearity hypothesis. ApEn begins to rise quickly

in the interval $3.55 < \mu \leq 3.70$, that contains the Feigenbaum attractor ($\mu \approx 3.57$)[2], although the comparison with the surrogates showed that the linearity hypothesis still passed for $3.55 \leq \mu \leq 3.58$. In the interval $3.58 < \mu < 3.60$ the linearity hypothesis finally began to be violated, and at $\mu = 3.60$, the linearity hypothesis was violated with $S > 50$. Thus this analysis identified the critical parameter to within one percent of the correct value.

2.7 Kaplan-Yorke map with varying parameter

The ApEn+FT analysis identified the transition between linear (i.e., periodic) and nonlinear behavior with the variation of a single parameter in the logistic map. We next examined the transition between chaotic to random behavior with one of the simplest of the Kaplan-Yorke family of maps[2], which is a two-dimensional extension of the logistic map for $\mu = 4$ (see Equation 7):

$$\begin{aligned} x_{n+1} &= 1 - 2x_n^2 \\ y_{n+1} &= ky_n + x_n \end{aligned} \tag{8}$$

This map has the following physical interpretation[2]: the y -coordinate describes the velocity of a unit mass particle in a dissipative medium subject to kicks of strength x_n at times $n\tau$. For a fixed dissipation rate γ , the dissipation parameter $k = \exp(-\gamma\tau)$. In the small- τ limit, $\sqrt{\tau}y_n$ describes a gaussian random process, with integral information dimension[2]. For moderate τ (i.e., $k \leq \frac{1}{2}$), however, Equation 8 describes complex chaotic dynamics, with a fractional information dimension. Figure 4 shows ApEn for the y -component of Equation 8, and its surrogates, for various k . Although the ApEn of the original series of y_n is flat for all but the highest values of k , the distance between the ApEn of the original and the surrogate data begins to smoothly decrease for $k > \frac{1}{2}$. The significance is at least 100 for $k \leq 0.95$. Finally, at $k = 0.99$, the ApEn of the original and the ApEn of the surrogate data begin to coincide ($S \approx 8$), consistent with an ARMA process. It is interesting that the value of ApEn alone did not indicate the transition from nonlinear to linear dynamics, since ApEn actually decreased for the random case. Instead, the transition became apparent only in conjunction with the surrogate data analysis.




Figure 3: ApEn and FT surrogates for the logistic map for various coupling constants. The circles and squares represent ApEn for the original and the surrogate series, respectively. The error bars represent the bootstrap estimate of the (0.5%, 99.5%) confidence interval for the surrogate data.




Figure 4: ApEn and surrogates for the Kaplan-Yorke map for various coupling constants. The solid circles show the ApEn of the original time series. The height of the error bars represents the bootstrap estimate of the (0.1%, 99.9%) confidence interval for the surrogate data. The dotted line (fit by a spline) represents the difference between the ApEn of the original data and the mean ApEn of the surrogate data.

3 Comments and Conclusion

The success of ApEn+FT analysis in detecting the nonlinearity in the noiseless chaotic models (i.e., Lorenz, Hénon, logistic, Kaplan-Yorke, Mackey-Glass) was gratifying, but was neither unique nor unexpected; indeed such success was necessary for further investigation. The example of the Mackey-Glass model also showed that the results could be quite sensitive to the length of the original time series. The ability of this analysis to distinguish between the complex chaotic behavior and the gaussian random behavior of the Kaplan-Yorke model, although not unique, adds to our confidence in the method. It was also encouraging to see that even relatively severe smoothing of the Hénon series was not sufficient to obscure its inherent nonlinearity. The ability of the ApEn+FT to detect nonlinear dynamics in the Santa Fe Institute Data Sets A (chaotic laser) and D (turbulence) data sets, although also not unexpected, was a more useful test because those data sets contained some noise, and in the increments considered here, were of intermediate length. The success of the linearity hypothesis for Data Set E with ApEn+FT was inconclusive, because this result might have been an accurate assessment of that data, or it might have been a failure to detect nonlinear dynamics in a noisy series. The strong sensitivity of this analysis to the noise in the Chua circuit suggests that, as with all other techniques[3][4][5], caution must be applied when the data is both noisy and short. Direct comparison to other methods was obstructed by the lack of equivalent analyses in the literature, and performing such analyses ourselves seemed beyond the scope of this paper. Nevertheless, our goal was to empirically assess the performance of the ApEn+FT method, and this has been obtained with the results of the last section.

The ApEn+FT method, like other implementations of surrogate data methods, cannot test for chaos, but only for the nonlinearity in the dynamics of a time series that, if present, merely indicates the potential for such a series also to be chaotic. The analysis can say nothing further concerning possible attractors, embedding dimension, or even if the data resulted from chaotic dynamics. Indeed one of the most spectacular violations of the linearity hypothesis came from the non-chaotic map of Grebogi, Ott, Pelikan and Yorke[38].

In addition to the problem of false positives (false success of the linearity hypothesis), there remains the possibility of false negatives (false violation of the linearity hypothesis) in ApEn+FT. For example, consider the sawtooth

series[1], with a spectrum that is poorly estimated by the straightforward Fast Fourier Transform employed in the FT method[29]. Each of the resulting surrogate time series would be much more irregular than the original time series, so that the ApEn of each would be significantly higher than the nearly zero ApEn of the original series, thereby yielding a false violation of the linearity hypothesis. However, to the extent that even complicated periodic systems might be easily recognized by visual inspection of either the series itself or its spectrum, the problem of false negatives might not pose the difficulties that false positives present. If more faithful surrogate time series were required, methods other than FT could be employed (see Introduction), and the ApEn method could be subsequently employed just as easily as we did here.

In conclusion, the ApEn statistic is easily implemented in the surrogate data technique to test the hypothesis of linear dynamics. Violations of the hypothesis were found with large significance for a wide variety of both experimental and computer-generated time series that were known or suspected to be governed by nonlinear dynamics. The ApEn+FT method appears to be useful especially for discriminating between stochastic but linear and complex but nonlinear time series of moderate size if the noise is not too large.

Acknowledgements

Steven Pincus provided us with valuable advice and reprints of his papers and presentations. Karen Leighly (Oklahoma), David Peak (Utah State), James Theiler (Los Alamos) and Jens Timmer (Freiburg) also provided valuable scientific correspondence.

A Algorithm for ApEn

The algorithmic interpretation of ApEn, i.e., the conditional probability that “...runs of patterns that are close for [some number] of observations remain close...” was developed by Pincus as follows[19]. For a sequence $\{u_k\}$, with $0 \leq k < N$, the $N - m$ window vectors $w(j)$ for $0 \leq j < N - m$, each of window length m , are defined by

$$w(j) \equiv [u_j, \dots, u_{j+m}]. \quad (9)$$

The fraction $C_i^m(r)$ of the distances from a given window vector $w(i)$ to all the window vectors (including itself) that lie within a tolerance r is defined by

$$C_i^m(r) \equiv \frac{1}{N - m} \sum_{j=0}^{N-m-1} \Theta(r - d[w(i), w(j)]), \quad (10)$$

where $\Theta(x \geq 0) = 1$ and $\Theta(x < 0) = 0$, and the distance d is given by the L_1 norm, i.e.,

$$d[w(i), w(j)] \equiv \max_{0 \leq k \leq m} |u_{i+k} - u_{j+k}|. \quad (11)$$

Then “the $C_i^m(r)$ measure, within a tolerance r , the regularity of patterns similar to a given pattern of window length m ”.[21] Collecting these ideas, Pincus defined ApEn by

$$\begin{aligned} \text{ApEn}(N, m, r) &\equiv \Phi^m(r) - \Phi^{m+1}(r), \\ \Phi^m(r) &\equiv \frac{1}{N - m} \sum_{i=0}^{N-m-1} \ln(C_i^m(r)) \end{aligned} \quad (12)$$

B ApEn and related quantities in various limits

There were at least two ways to employ ApEn to measure regularity in data, but only one of these was fruitful for our application. The differences arise

from the way the parameters m , r , and N employed. Viewed in the limits $N \rightarrow \infty$, $m \rightarrow \infty$, and $r \rightarrow 0$, Pincus showed that the definitions above yield, at least theoretically[19], the more familiar correlation integral $\Gamma_m(r)$ and correlation dimension β as follows:

$$\Gamma_m(r) = \lim_{N \rightarrow \infty} \frac{1}{N - m} \sum_{i=0}^{N-m-1} C_i^m(r), \quad (13)$$

$$\beta = \lim_{\substack{m \rightarrow \infty \\ r \rightarrow 0}} \frac{\ln(\Gamma_m(r))}{\ln(r)}. \quad (14)$$

The theoretical bounds to ApEn in these limits are 0 for a perfectly regular sequence, and $\ln(B)$ for a maximally irregular (random) sequence of base B numbers. ApEn itself theoretically converges to the Eckmann-Ruelle entropy[19][20][22]. Pincus has explored the large- N limit numerically for binary numbers, taking advantage of their special structure[23]. However, for most work with decimal numbers, the convergence is so slow that the limits above are practically unobtainable, therefore we did not attempt to employ them in this work.

C Connection between ApEn of binary strings and the Ising model

As we show in Figure 5, ApEn, even with $m = 2$, essentially reproduces the thermodynamic entropy of a one-dimensional zero-field Ising model, a two-state spin model in which the ground state is anti-ferromagnetic, i.e., $\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow \dots$, with nearest-neighbor interactions on a periodic chain[2][43][44]. This coincidence seems not to have been noticed before, and because it provides a physical analogy for ApEn, we digress here briefly to include a few details of a calculation we carried out. Although the thermodynamics (including the entropy) are known exactly for this model[43], in order to apply ApEn, we generated the equilibrium configurations (a string of 1024 bits) from a Monte-Carlo simulation that employed importance sampling and periodic boundary conditions[44]. For each temperature we generated a total of 50 000 equilibrium configurations (binary strings), and computed the mean $\text{ApEn}(N = 1024, m = 2, r < 1)$ from 20 configurations that were separated by 2000 passes, after discarding the first 10 000 configurations. The comparison of the mean ApEn with the exact thermodynamic entropy is displayed

in the Figure 5. We estimated the standard deviation of the mean ApEn to be between 0.001 and 0.002 for all temperatures. The high temperature limit corresponds to $\ln(2)$. There is a rough correspondence between the behavior of ApEn with varying temperature T for the Ising model and as a function of the dissipation rate γ for the Brownian noise model (see Figure 1).




Figure 5: ApEn and the exact entropy for equilibrium configurations of the Ising model for various inverse temperatures. The solid line is the exact thermodynamic entropy, and the circles are the calculated mean ApEn. The estimated standard deviations are much smaller than the circle diameter.

References

- [1] G. L. Baker and J. P. Gollub, *Chaotic Dynamics*, (Cambridge University Press, New York, 1994).
- [2] C. Beck and F. Schlogl, *Thermodynamics of chaotic systems*, (Cambridge University Press, New York, 1995).
- [3] D. Kaplan and L. Glass, *Understanding Nonlinear Dynamics*, (Springer-Verlag, New York, 1995).
- [4] H. D. I. Abarbanel, *Analysis of Observed Chaotic Data*, (Springer-Verlag, New York, 1996).
- [5] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*, (Cambridge University Press, New York, 1997).
- [6] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer, “Testing for nonlinearity in time series: the method of surrogate data”, *Physica D* **58** (1992) 77; reprinted in *Coping with Chaos*, edited by E. Ott, T. Sauer, and J. A. Yorke (John Wiley, New York, 1994).
- [7] M. B. Kennel and S. Isabelle, “Method to distinguish possible chaos from colored noise and to determine embedding parameters”, *Phys. Rev. A* **46** (1992) 3111.
- [8] J. Theiler and D. Prichard, “Constrained-Realization Monte-Carlo method for Hypothesis Testing”, *Physica D* **94** (1996) 221.
- [9] C. J. Stam, B. Jelles, H. A. M. Achtereekte, S. A. R. B. Rombouts, J. P. J. Slaets, and R. W. M. Keunen, “Investigation of EEG nonlinearity in dementia and Parkinson’s disease”, *Electroen. Clin. Neuro.* **95** (1995) 309.
- [10] C. Diks, J. C. van Houwelingen, F. Takens, J. DeGoede, “Reversibility as a criterion for discriminating time series”, *Phys. Lett. A* **201** (1995) 221.
- [11] M. Palus, “Testing for nonlinearity using redundancies: quantitative and qualitative aspects”, *Physica D* **80** (1995) 186.

- [12] C. Schittenkopf and G. Deco, “Testing nonlinear Markovian hypotheses in dynamical systems”, *Physica D* **104** (1997) 61.
- [13] L. Kocarev and D. M. Walker, “Compactness of symbolic sequences from chaotic systems”, *Phys. Lett. A* **274** (2000) 200.
- [14] J. Timmer, “The power of surrogate data testing with respect to non-stationarity”, *Phys. Rev. E* **58** (1998) 5153.
- [15] T. Schreiber, “Interdisciplinary application of nonlinear time series method”, *Phys. Rep.* **308** (1999) 1.
- [16] D. Kugiumtzis, “Test your surrogate data before you test for nonlinearity”, *Phys. Rev. E* **60** (1999) 2808; D. Kugiumtzis, “Surrogate data test for nonlinearity including nonmonotonic transforms”, *ibid.* **62** (2000) R25.
- [17] J. Theiler, “Spurious dimension from correlation algorithms applied to limited time-series data”, *Phys. Rev. A* **34** (1986) 2427.
- [18] J.-P. Eckmann and D. Ruelle, “Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems”, *Physica D* **56** (1992) 185; reprinted in *Coping with Chaos*, edited by E. Ott, T. Sauer, and J. A. Yorke (John Wiley, New York, 1994).
- [19] S. M. Pincus, “Approximate entropy as a measure of system complexity”, *Proc. Natl. Acad. Sci. USA* **88** (1991) 2297.
- [20] S. M. Pincus and W.-M. Huang, “Approximate entropy: Properties and Applications”, *Commun. Statist.-Theory Meth.* **21** (1992) 3061.
- [21] S. M. Pincus, “Quantifying Complexity and Regularity of Neurobiological Systems”, in *Methods in Neurosciences*, Vol. 28 (Academic Press, Orlando, 1995).
- [22] S. M. Pincus, “Approximate entropy (ApEn) as a complexity measure”, *Chaos* **5** (1995) 110.
- [23] S. M. Pincus and R. E. Kalman, “Not all (possibly) random sequences are created equal”, *Proc. Natl. Acad. Sci. USA* **94** (1997) 3513.

- [24] S. Chatterjee, M. R. Yilmaz, M. Habibullah, M. Laudato, “An Approximate Entropy Test for Randomness”, *Commun. Statist.-Theory Meth.* **29** (2000) 655.
- [25] S. M. Pincus and A. L. Goldberger, “Physiological time-series analysis: what does regularity quantify?”, *Am. J. Physiol.* **266** (1994) H1643.
- [26] L. J. Groome, D. M. Mooney, S. B. Holland, L. A. Smith, J. L. Atterbury, and P. C. Loizou, “Human fetuses have nonlinear cardiac dynamics”, *J. Appl. Physiol.* **87** (1999) 530.
- [27] D. Cysarz, H. Bettermann, P. Van Leeuwen, “Entropies of short binary sequences in heart period dynamics”, *Am. J. Physiol.* **278** (2000) H2163.
- [28] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap* (Chapman and Hall, New York, 1993).
- [29] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, Second edition, (Cambridge University Press, New York, 1994); see Chapters 12 and 13.
- [30] D. A. Dahl, C. L. Atwood, and R. A. LaViolette, “A random walk pseudorandom byte generator”, *Appl. Math. Modelling* **24** (2000) 771.
- [31] S. O. Rice, “Mathematical analysis of random noise, pt. II”, in *Selected Papers on Noise and Stochastic Processes*, edited by N. Wax (Dover, New York, 1954).
- [32] A. R. Osborne and A. Provenzale, “Finite correlation dimension for stochastic systems with power-law spectra”, *Physica D* **35** (1989) 357.
- [33] P. Gaspard and X.-J. Wang, “Noise, chaos, and (ϵ, τ) -entropy per unit time”, *Phys. Rep.* **235** (1993) 291.
- [34] A. S. Weigend and N. A. Gershenfeld, *Time Series Prediction: Forecasting the Future and Understanding the Past*, (Addison-Wesley, Reading, 1994)

- [35] The data employed in the competition are described by Weigend and Gershenfeld[34]. The data are publically available at the Santa Fe Institute ftp site at <ftp://ftp.santafe.edu/pub/Time-Series>. We downloaded the data, along with detailed descriptions and references, from Weigend’s web site at <http://www.stern.nyu.edu/~aweigen/Time-Series/SantaFe.html>
- [36] L. A. Aguirre, G. G. Rodrigues, and E. M. Mendes, “Nonlinear identification and cluster analysis of chaotic attractors from a real implementation of Chua’s circuit”, *Int. J. of Bifurcation and Chaos* **7** (1997) 1411.
- [37] The data set DSVC1.dat, labeled “chua” here, was downloaded from <http://www.cpdee.ufmg.br/~MACSIN/services/data/data.htm>. That data set is represented as relatively noise-free. Other data sets SPVC1.dat and DSIL.dat are longer but are represented as noisy, although we could not determine the extent of the noise quantitatively. ApEn+FT analysis passes the linearity hypothesis for the latter two data sets, suggesting that the noise was enough to obscure the nonlinear dynamics.
- [38] C. Grebogi, E. Ott, S. Pelikan, and J. A. Yorke, “Strange attractors that are not chaotic”, *Physica D* **13** (1984) 261; see Equation 4.
- [39] We employed 6400 bootstrap replicates of the 10-element set of ApEns. In all cases the mean and the median practically coincided, so that we did not need nor employ the corrections generated by more elaborate (e.g., BC_A) methods generally recommended by Efron and Tibshirani[28].
- [40] K. M. Leighly and P. T. O’Brien, “Evidence for nonlinear X-ray variability from the broad-line radio galaxy 3C 390.3”, *Astrophys. J.* **481** (1997) L15.
- [41] J. Timmer, U. Schwarz, H. U. Voss, I. Wardinski, T. Belloni, G. Hasinger, M. van der Klis, and J. Kurths, “Linear and nonlinear time series analysis of the black hole candidate Cygnus X-1”, *Phys. Rev. E* **61** (2000) 1342.
- [42] J. Theiler and S. Eubank, “Don’t bleach chaotic data”, *Chaos* **3** (1993) 771.

- [43] C. J. Thompson, Mathematical Statistical Mechanics, (Princeton University Press, Princeton, 1972).
- [44] D. Chandler, Introduction to Modern Statistical Mechanics, (Oxford University Press, New York, 1987).